

Intel[®] Omni-Path Fabric Software

Release Notes for 10.3

January 2017



You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting: <http://www.intel.com/design/literature.htm>

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at <http://www.intel.com/> or from the OEM or retailer.

Intel, Intel Xeon Phi, Xeon, and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.


*Other names and brands may be claimed as the property of others.

Copyright © 2015-2017, Intel Corporation. All rights reserved.



Contents

1.0 Overview of the Release	4
1.1 Introduction.....	4
1.2 Audience	4
1.3 Software License Agreement	4
1.4 If You Need Help	4
1.5 New Enhancements and Features in this Release	4
1.6 Supported Features	5
1.7 Release Packages	5
1.8 Firmware Files	6
1.9 Operating Systems	6
1.10 Parallel File Systems	7
1.11 Compilers.....	7
1.11.1 MPI	7
1.11.2 MVAPICH2 and Open MPI	7
1.12 Hardware	8
1.13 Installation Requirements	8
1.13.1 Software and Firmware Requirements.....	8
1.13.2 Installation Instructions	8
1.14 Product Constraints	9
1.15 Product Limitations	9
1.16 Documentation.....	9
2.0 Issues	11
2.1 Introduction.....	11
2.2 Resolved Issues	11
2.3 Open Issues	12
Tables	
1-1 Firmware Files.....	6
1-2 Operating Systems Supported	6
1-3 MPI Compilers	7
1-4 MVAPICH2 and Open MPI	7
1-5 Hardware Supported	8
1-6 Related Documentation for this Release	9
2-1 Resolved Issues.....	11
2-2 Open Issues.....	12



1.0 Overview of the Release

1.1 Introduction

This document provides a brief overview of the changes introduced into the Intel® Omni-Path Software by this release. References to more detailed information are provided where necessary. The information contained in this document is intended as supplemental information only; it should be used in conjunction with the documentation provided for each component.

These Release Notes list the features supported in this software release, open issues, and issues that were resolved during release development.

1.2 Audience

The information provided in this document is intended for installers, software support engineers, service personnel, and system administrators.

1.3 Software License Agreement

This software is provided under license agreements and may contain third-party software under separate third-party licensing. Please refer to the license files provided with the software for specific details.

1.4 If You Need Help

Technical support for Intel® Omni-Path products is available 24 hours a day, 365 days a year. Please contact Intel Customer Support or visit www.intel.com for additional detail.

1.5 New Enhancements and Features in this Release

The following enhancements and features are new for the 10.3 release:

- Additional operating systems supported. See [Table 1-2](#) for details.
- Added support for Enterprise Edition for Lustre* software version 3.1.
- Support for the Enhanced Hypercube Routing Engine is outside the scope of Intel® OPA support. However, Intel partners may offer such support as part of their solutions. In addition there is an open source community who may be able to answer specific questions and provide guidance with respect to the Enhanced Hypercube Routing Engine.
- Limited support for active optical cables (AOC) on server platforms using integrated HFI for OPA (commonly known as "-F").



1.6 Supported Features

- The list of supported operating systems is in [Table 1-2](#).
- The list of supported hardware is in [Table 1-5](#).
- Coexistence with Intel® True Scale Architecture. This release supports True Scale hardware serving as an InfiniBand* storage network with the Intel® Omni-Path hardware used for computing. Note that connecting a True Scale adapter card to an Omni-Path switch, or vice-versa, is not supported. For more details on this feature, refer to *Intel® Omni-Path Fabric Host Software User Guide*.
- Supports Dual Rail: Two Intel® Omni-Path Host Fabric Interface (HFI) cards in the same server connected to the same fabric
- Supports Dual Plane: Two HFI cards in the same server connected to separate fabrics.
- Limited validation testing performed on network storage file systems:
 - NFS over TCP/IP
- Active Optical Cables. For details, see the Cable Matrix at: <http://www.intel.com/content/www/us/en/high-performance-computing-fabrics/omni-path-cables.html>
- MPI applications are provided in a stand-alone rpm package.
- Intel® Xeon® v4 processor (codename Broadwell) support
- Intel® Xeon Phi™ support
- Monitored Intel® Omni-Path Host Fabric Interface
- DHCP and LDAP supported on Intel® Omni-Path Edge Switch 100 Series and Intel® Omni-Path Director Class Switch 100 Series hardware.

1.7 Release Packages

There are two Intel® Omni-Path Fabric Software packages:

- Basic for compute nodes
- IFS for the management node

The Basic package includes:

- Software that installs the following packages to the distribution OpenFabrics Alliance* (OFA):
 - libibumad is based on the RHEL* or SLES* release package. It contains Intel patches that support Intel® Omni-Path Architecture (Intel® OPA) technology.
 - ibacm is the latest upstream code applied with RHEL* patches.
 - hfi1-firmware, hfi1-psm, hfi1-diagtools-sw, libhfi1verbs
 - Open MPI built for verbs and PSM2 using gcc, and Intel compilers.
 - MVAPICH2 built for verbs and PSM2 using gcc, and Intel compilers.
 - mpitests
 - mpi-selector
 - GASnet
 - openSHMEM
 - srptools (includes the latest upstream code)
 - Firmware files listed in [Table 1-1](#).



- compat-rdma which delivers kernel changes based on the OFA version. The components installed are the hfi1 driver and Intel-enhanced versions of other kernel packages. See the *Building Lustre* Servers with Intel® Omni-Path Architecture Application Note* for details.

Note: In the Intel® Omni-Path Software package for RHEL* 7.2, the hfi1 driver and ifs-kernel-updates are supplied as a smaller package.

The IFS package includes the Basic package plus:

- Fabric Manager, which allows comprehensive control of administrative functions using a mature Subnet Manager. Fabric Manager simplifies subnet, fabric, and individual component management, easing the deployment and optimization of large fabrics.
- Fabric Suite FastFabric Toolset, which enables rapid, error-free installation and configuration of Intel® OPA host software and management software tools, as well as simplified installation, configuration, validation, and optimization of HPC fabrics. For details, refer to the Fabric Suite FastFabric documentation in [Table 1-6](#).

1.8 Firmware Files

This release of the Intel® Omni-Path Software contains the firmware files listed in [Table 1-1](#).

Table 1-1. Firmware Files

Description	File Name	Version
HF11 UEFI Option ROM	HfiPcieGen3_1.3.0.0.0.efi	1.3.0.0.0
UEFI UNDI	HfiPcieGen3Loader_1.3.0.0.0.rom	1.3.0.0.0
HF11 SMBus Microcontroller Firmware (Thermal Monitor)	hfi1_smbus.fw	10.2.1.0.3

1.9 Operating Systems

This release of the Intel® Omni-Path Software supports the operating systems listed in [Table 1-2](#).

Table 1-2. Operating Systems Supported

Operating System	Update/SP	Kernel Version
Red Hat* Enterprise Linux* (RHEL*) 7.2 X86_64	Update 2	3.10.0-327.el7.x86_64
Red Hat* Enterprise Linux* (RHEL*) 7.3 X86_64	Update 3	3.10.0-513.el7.x86_64
CentOS* 7.2 X86_64	N/A	3.10.0-327.el7.x86_64
Scientific Linux* 7.2 X86_64	N/A	3.10.0-327.el7.x86_64
SUSE* Linux* Enterprise Server (SLES*) 12.1 X86_64	Service Pack 1	3.12.49-11.1-default
SUSE* Linux* Enterprise Server (SLES*) 12.2 X86_64	Service Pack 2	4.4.21-69-default



1.10 Parallel File Systems

The following parallel file systems have been tested with this release of the Intel® Omni-Path Software:

- Intel® Enterprise Edition Lustre* (IEEL) 3.1
 - RHEL* versions supported by Intel® Omni-Path Software.
- IBM* General Parallel File System (GPFS) version 4.0.1
 - RHEL* 7.2.

Refer to the *Intel® Omni-Path Fabric Performance Tuning User Guide* for details on optimizing parallel file system performance with Intel® Omni-Path Software.

1.11 Compilers

1.11.1 MPI

This release supports the following MPI implementations:

Table 1-3. MPI Compilers

MPI Implementation	Runs Over	Compiled With
Open MPI 1.10.4	Verbs	GCC
	PSM2	GCC, Intel
MVAPICH2-2.1	Verbs	GCC
	PSM2	GCC, Intel
Intel® MPI 5.1.3	Verbs	GCC
	PSM2	GCC, Intel


1.11.2 MVAPICH2 and Open MPI

MVAPICH2 and Open MPI have been compiled for PSM2 to support the following versions of the compilers:

Table 1-4. MVAPICH2 and Open MPI

Compiler	Linux* Distribution	Compiler Version
(GNU) gcc	RHEL* 7.2	gcc (GCC) 4.8.5 20150623 (Red Hat* 4.8.5-4)
(GNU) gcc	RHEL* 7.3	gcc (GCC) 4.8.5 20150623 (Red Hat* 4.8.5-11)
(GNU) gcc	SLES* 12 SP 1	gcc (SUSE* Linux*) version 4.8.5
(GNU) gcc	SLES* 12 SP 2	gcc (SUSE* Linux*) version 4.8.5
(Intel) icc	RHEL* 7.2	icc (ICC) 15.0.1
(Intel) icc	RHEL* 7.3	icc (ICC) 15.0.1
(Intel) icc	SLES* 12 SP 1	icc (ICC) 15.0.1
(Intel) icc	SLES* 12 SP 2	icc (ICC) 15.0.1

Note: Refer to the *Intel® Omni-Path Fabric Host Software User Guide* for set up information when using Open MPI with the SLURM PMI launcher and PSM2.



1.12 Hardware

Table 1-5 lists the hardware supported in this release.

Table 1-5. Hardware Supported

Hardware	Description
Intel® Xeon® v3 Processor	Haswell CPU-based servers
Intel® Xeon® v4 Processor	Broadwell CPU-based servers
Intel® Xeon Phi™ Processor	Knights Landing CPU-based servers
Intel® Omni-Path Host Fabric Interface 100HFA016 (x16)	Single Port Host Fabric Interface (HFI)
Intel® Omni-Path Host Fabric Interface 100HFA018 (x8)	Single Port Host Fabric Interface (HFI)
Intel® Omni-Path Switch 100SWE48Q	Managed 48-port Edge Switch
Intel® Omni-Path Switch 100SWE48U	Externally-managed 48-port Edge Switch
Intel® Omni-Path Switch 100SWE24Q	Managed 24-port Edge Switch
Intel® Omni-Path Switch 100SWE24U	Externally-managed 24-port Edge Switch
Intel® Omni-Path Director Class Switch 100SWD24	Director Class Switch 100 Series, up to 768 ports
Intel® Omni-Path Director Class Switch 100SWD06	Director Class Switch 100 Series, up to 192 ports

1.13 Installation Requirements

1.13.1 Software and Firmware Requirements

Table 1-2 lists the operating systems supported by this release. Refer to the *Intel® Omni-Path Fabric Software Installation Guide* for the required packages.

1.13.2 Installation Instructions

There are two Intel® Omni-Path Fabric Software packages:

- IntelOPA-IFS.<distro>-x86_64.<version>.tgz for the management node.
- IntelOPA-Basic.<distro>-x86_64.<version>.tgz for compute nodes.

The packages in the tgz file are RPMs. Installing individual RPMs is not supported in the 10.3 release.

Refer to the *Intel® Omni-Path Fabric Software Installation Guide* for related software requirements and complete installation procedures. Refer to the *Intel® Omni-Path Fabric Hardware Installation Guide* for related firmware requirements.

1.13.2.1 Required Pre-Installation to perform external modules builds on SLES* 12 Systems

SLES* 12 kernel-development environment is not ready "out-of-box" for external modules build. It has to be prepared prior to the installation. To rebuild SLES* 12 kernel pieces, perform the following steps:

1. Change directory:

```
cd /lib/modules/3.12.28-*****/source
```




2. Create the following files:
 - Make cloneconfig
 - Make oldconfig
 - Make __headers
3. Only build as needed.

1.14 Product Constraints

None.

1.15 Product Limitations

This release has the following product limitations:

- The embedded version of the Fabric Manager supports a maximum of 100 nodes within a fabric. This is due to the limited memory and processing resources available in the embedded environment.
- PA Failover should **not** be enabled with FMs running on differing software versions. PA Failover is enabled via configuration: `<PM>/<ImageUpdateInterval> > 0`
- Enabling UEFI Optimized Boot on some platforms can prevent the HFI UEFI driver from loading during boot. To prevent this, do not enable UEFI Optimized Boot.

1.16 Documentation

Table 1-6 lists the end user documentation for the current release.

Documents are available at the following URLs:

- Intel® Omni-Path Switches Installation, User, and Reference Guides
www.intel.com/omnipath/SwitchPublications
- Intel® Omni-Path Fabric Software Installation, User, and Reference Guides
www.intel.com/omnipath/FabricSoftwarePublications
- Drivers and Software (including Release Notes)
www.intel.com/omnipath/downloads


Table 1-6. Related Documentation for this Release (Sheet 1 of 2)

Document Title
Hardware Documents
<i>Intel® Omni-Path Fabric Switches Hardware Installation Guide</i>
<i>Intel® Omni-Path Fabric Switches GUI User Guide</i>
<i>Intel® Omni-Path Fabric Switches Command Line Interface Reference Guide</i>
<i>Intel® Omni-Path Edge Switch Platform Configuration Reference Guide</i>
<i>Intel® Omni-Path Fabric Switches Release Notes (includes managed and externally-managed switches)</i>
<i>Intel® Omni-Path Host Fabric Interface Installation Guide</i>
Fabric Software Documents
<i>Intel® Omni-Path Fabric Software Installation Guide</i>



Table 1-6. Related Documentation for this Release (Sheet 2 of 2)

Document Title
<i>Intel® Omni-Path Fabric Suite Fabric Manager User Guide</i>
<i>Intel® Omni-Path Fabric Suite FastFabric User Guide</i>
<i>Intel® Omni-Path Fabric Host Software User Guide</i>
<i>Intel® Omni-Path Fabric Suite Fabric Manager GUI Online Help</i>
<i>Intel® Omni-Path Fabric Suite Fabric Manager GUI User Guide</i>
<i>Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide</i>
<i>Intel® Performance Scaled Messaging 2 (PSM2) Programmer's Guide</i>
<i>Intel® Omni-Path Fabric Performance Tuning User Guide</i>
<i>Intel® Omni-Path Host Fabric Interface Platform Configuration Reference Guide</i>
<i>Intel® Omni-Path Fabric Software Release Notes</i>
<i>Intel® Omni-Path Fabric Manager GUI Release Notes</i>
<i>Intel® Omni-Path Storage Router Design Guide</i>
<i>Intel® Omni-Path Fabric Staging Guide</i>
<i>Building Lustre* Servers with Intel® Omni-Path Architecture Application Note</i>



2.0 Issues

2.1 Introduction


This section provides a list of the resolved and open issues in the Intel® Omni-Path Software.

2.2 Resolved Issues

Table 2-1 lists issues that are resolved in this release.

Table 2-1. Resolved Issues

ID	Component	Description	Resolved in Release
133377	HFI Host Driver	irqbalance settings are not being honored correctly after a reboot.	10.3
133707	Software Installation	Updating to the RHEL* 7.2 kernel for the CVE-2016-0728 update in OSes prior to 7.2 causes the Intel® Omni-Path installation to fail.	10.3
134111	Software Configuration Management	On some older HFI and HFI-like cards, running hfi1_eprom -V -c to inquire the version of the AOC configuration file on the card may return an invalid version of "etnIRFWI".	10.3
134124	Fabric Management Software/SM	HFI port stuck in INIT state due to SM failure to set pkeys.	10.3
134135 134429	HFI Host Fabric Software/ HFI Host Driver	When running communication-intensive workloads with 10KB MTU, it is possible to encounter node and/or job failures.	10.3
134283	Software Installation/ Packaging	When downgrading on a SLES* 12.X system from Intel® OPA version 10.2.X to a previous version, the following install errors occur: ERROR - Failed to install and error: Failed dependencies: libibmad5 is needed by opa-basic-tools...	10.3
134772	Fabric Management Tools/Basic	opatmmtool will fail if provided with a filename (full path) that is longer than 63 characters.	10.3
135000	Fabric Management Software	Fabric Manager configuration files that specify IncludeGroup fields with undefined or nonexistent device groups could cause Fabric Manager failure.	10.3
136318	Fabric Management Software	SM crashes showing segfault errors in logs and high CPU usage. These crashes were caused by a mismatch of pahistory file versions.	10.3



2.3 Open Issues

Table 2-2 lists the open issues for this release.

Table 2-2. Open Issues (Sheet 1 of 6)

ID	Component	Description	Workaround
129563	HFI Host Fabric Software/ MPI	Memory allocation errors with Mvapich2-2.1/Verbs.	When running mvapich2 jobs with a large number of ranks (for example, between 36 and 72 ranks), you must set the following parameters in <code>/etc/security/limits.conf</code> : * hard memlock unlimited * soft memlock unlimited Also, you must increase the <code>lkey_table_size</code> : LKEY table size in bits (2^n , $1 \leq n \leq 23$) from its default of 16 to 17. For instructions on setting module parameters, refer to Appendix A in the <i>Intel® Omni-Path Fabric Performance Tuning User Guide</i> .
130336	HFI Host Fabric Software/ Tools	<code>hfi1stats</code> cannot be run at user level due to mount-point privileges.	The administrator can provide <code>sudo</code> access to <code>hfi1stats</code> or provide root access to users.
131017	Open Software	Verbs <code>ib_send_bw</code> , <code>ib_read_bw</code> , and <code>ib_write_bw</code> are not working with the <code>-R</code> option to use the RDMA CM API to create QPs and exch data.	The following combinations of client and server DO NOT allow RDMA CM connections: Client Server SLES* 12.1 SLES* 12.1 or (intermittent) RHEL* 7.2 Using SLES* clients results in "Unexpected CM event bl blka 0" errors. Additionally, there are long (5-10 sec) initial delays when using these combinations: Client Server SLES* 12.1 RHEL* 7.2 SLES* 12.0 RHEL* 7.2
131745	HFI Host Fabric Software/ MPI	When running OpenMPI 1.10.0 on SLES* 12 with large number of ranks per node (over 40), it may happen that the ORTE daemon (orted) "hangs" during the finalization of job.	Stopping and resuming the "hung" orted process allows the job to finish normally. To find the hung process, run the <code>ps</code> and find a node with several job zombie processes. In that same node, identify the orted process ID and send a stop signal (kill -19 <PID>) and a continue signal (kill -18 <PID>).
132207	Open Software	Kernel crash caused by the <code>ib_srpt</code> module.	Install this kernel patch: https://git.kernel.org/cgit/linux/kernel/git/torvalds/linux.git/commit/?id=51093254bf879bc9ce96590400a87897c7498463
132219	HFI Host Fabric Software	Server platforms running IFS 10.3.0 release (or Intel® OPA software delivered in certain Linux* OS distributions) and using integrated HFI for OPA (commonly known as "-F") may not support Active Optical Cables (AOC) after boot up.	Contact your server vendor for the appropriate <code>hfi1_platform.dat</code> file and install it in the <code>lib/firmware/updates</code> folder.
133596	Software Installation/ Packaging	The install script does not check for packages that are necessary prerequisites for installation.	Refer to the <i>Intel® Omni-Path Fabric Software Installation Guide</i> , in the OS RPMs Installation Prerequisites section for the complete list.
133604	Open Software	Bonding driver incorrectly shows hardware address of IPoIB interfaces.	Use the <code>opainfo</code> command to retrieve the PortGUID and <code>ip addr show ib0</code> to get the correct 20-byte hardware address of OPA network interface.



Table 2-2. Open Issues (Sheet 2 of 6)

ID	Component	Description	Workaround
133633	Open Software	OpenMPI and Mvapich2 compiles fail to link properly when using the Intel compilers.	No workaround available.
134268	HFI Host Pre-boot Software	The Option ROM image (e.g. containing a UEFI driver) may not be executed if the BIOS configures the HFI Expansion ROM BAR with an address that is not 16MB aligned.	Use a BIOS that configures the HFI Expansion ROM BAR with an address that is 16MB aligned. Note that the main memory BAR of the device is 64MB in size and therefore requires 64MB alignment. A BIOS implementation that places the Expansion ROM BAR immediately after the main memory BAR automatically provides this workaround. Many BIOS implementations have this property and automatically meet the workaround criteria.
134353	Link Working Group Issues	Very infrequently, when a link goes down, the logical link state can remain stuck in the 'Init' state.	The device containing the affected port must be rebooted in order to resolve the issue. Ports in this state have a logical link state of 'Init' but do NOT have a physical port state of 'LinkUp'.
134471	HFI Host Pre-boot Software	The HFI UEFI driver cannot boot via PXE using Grub 2.	Use Elilo instead.
134493	HFI Host Fabric Software/ MPI	When using Mvapich2 with Intel® Omni-Path PSM2, users will notice unexpected behavior when seeding the built-in random number generator with functions like srand or srandom before MPI_Init is called. MPI_Init re-seeds the random number generator with its own value and does not restore the seed set by the user application. This causes different MPI ranks to generate different sequences of random numbers even though they started with the same seed value.	Seed the random number generator after MPI_Init is called or use the reentrant random number generator functions such as drand48_r.
134494	HFI Host Fabric Software/ MPI	Open MPI uses srand() family functions at MPI_Init() time. Therefore, if the user sets srand() before calling MPI_Init(), the values will be altered.	a) Fixed in Open MPI 2.0.1. b) Call srand() functions family after calling MPI_Init().
134819	HFI Host Pre-boot Software	In KNL-F EFI shell, the command ifconfig -l does not correctly display the IP address after being assigned via DHCP.	Launch a newer, working version of the EFI shell from the embedded shell.
134821	HFI Host Pre-boot Software	The UEFI network stack is initialized with a default network address before the driver receives a MAD packet containing an updated and actual subnet prefix. Therefore, in ARP and IP UEFI drivers the old (default) HW address is still used, causing problems with packet receiving and transmitting.	Use default subnet prefix, 0xfe80000000000000, when configuring a subnet for PXE boot over OPA.
134866	Fabric Management Software	hostverify.sh cannot properly detect if SRP is enabled on target node.	Modify hostverify.sh with the following: Replace autostart=\$(systemctl is-enabled srpd 2>&1 grep enabled) with autostart=\$(lsmod grep ib_srp)
134904	Custom HFI SW/FW	Legacy PXE boot using iPXE while the HFI UEFI driver is loaded causes a hang.	Configure PXE operation to boot using UEFI boot mode.
134956	Fabric Management Tools/FastFabric	ib0 fails to become ready on warm reboots.	If the HFI port transition to ACTIVE state takes an unusually long time and ipoib interface remains down, verify that the NM_CONTROLLED=no option is not present in the ifcfg-ib* files.



Table 2-2. Open Issues (Sheet 3 of 6)

ID	Component	Description	Workaround
135040	Fabric Management Software	You can't currently specify portions of an Intel® DCS chassis that is not populated and is not expected to be populated. If CoreFull is 1, all the internal links for that chassis are generated when run against <code>opaxlatopology</code> . If CoreFull is 0, none of the links are generated.	Copy internal configuration of desired Intel® DCS switch into the main topology tab of the spreadsheet. Then delete all lines corresponding to leafs or spines that are not present in the configuration.
135068	HFI Host Pre-boot Software	When PXE booting using older versions of Grub 2 over Ethernet while the HFI UEFI driver is loaded, some servers will crash with an RSOD (Red Screen of Death).	Upgrade to the latest version of Grub 2.
135084	HFI Host Hardware	During extensive power cycle testing, HFI adapter might fail to appear in PCI config space.	Reboot or power cycle the platform.
135180	HFI Host Driver	OpenMPI/PSM2 timeouts during MPI stress tests on Haswell and Intel® Xeon® mixed fabrics.	Load the hfi1 module with the following parameter: <code>sudo modprobe hfi1 rcvhdrcnt=16352</code>
135259	DC Link Software	On rare occasions, the HFI links do not come up after a reboot.	Reboot or bounce the link.
135326	HFI Host Pre-boot Software	Calling <code>opasmaquery</code> fails when called from a non-SM node to a node which has not booted to the OS.	Use the SM node when calling <code>opasmaquery</code> in this way.
135355	Software Installation/Packaging	Due to changes in where the IFS packages are installed, customers using the FastFabric tools and upgrading to 10.3 from an earlier release must find each occurrence of <code>/opt/opa</code> in the <code>opafastfabric.conf</code> file and replace the string with <code>/usr/lib/opa</code> .	This includes: <pre>export FF_MPI_APPS_DIR=\${FF_MPI_APPS_DIR:-/opt/opa/src/mpi_apps} export FF_SHMEM_APPS_DIR=\${FF_SHMEM_APPS_DIR:-/opt/opa/src/shmem_apps} export FF_PRODUCT=\${FF_PRODUCT:-IntelOPA-Basic.`cat /opt/opa/tools/osid_wrapper`}</pre> These lines should be changed to: <pre>export FF_MPI_APPS_DIR=\${FF_MPI_APPS_DIR:-/usr/lib/opa/src/mpi_apps} export FF_SHMEM_APPS_DIR=\${FF_SHMEM_APPS_DIR:-/usr/lib/opa/src/shmem_apps} export FF_PRODUCT=\${FF_PRODUCT:-IntelOPA-Basic.`cat /usr/lib/opa/tools/osid_wrapper`}</pre>



Table 2-2. Open Issues (Sheet 4 of 6)

ID	Component	Description	Workaround
135390	HFI Host Driver	<p>For certain older HFI adapters, the card may not be flashed with the AOC configuration file, or may be flashed with an older version of the AOC configuration file. With release 10.3, these adapters will fail to link up with AOCs, and these messages may be seen in dmesg:</p> <pre>[26.903186] hfi1 0000:d5:00.0: hfi1_0: parse_platform_config:Bad config file [26.903186] hfi1 0000:d5:00.0: hfi1_0: parse_platform_config:File claims to be larger than read size In addition to: [27.351555] hfi1 0000:d5:00.0: hfi1_0: tune_serdes: Unknown port type</pre>	<p>For HFI and HFI-equivalent customers, the resolution for this issue is to flash the AOC configuration file provided in <code>/lib/firmware/updates/hfi1_platform.dat</code> to the card using the following command line:</p> <pre># hfi1_eprom -w -c /lib/firmware/updates/hfi1_platf orm.dat -o /opt/opa/bios_images/HfiUndiStub _<version>.rom -b /opt/opa/bios_images/HfiPcieGen3 _<version>.efi -d /sys/bus/pci/devices/<PCI address>/resource0</pre>
135545	Fabric Management Software	A change has been made to several SA record attributes which causes incompatibilities between the Fabric tool suite and the SA.	You must update both the SA and the Fabric tool suite at the same time to avoid this incompatibility.
135648	Fabric Management Tools/FastFabric	MPI applications are installed under the <code>/usr/lib</code> directory structure, which may be set up to be read-only overall. This causes resulting FastFabric operations to fail since <code>mpi_apps</code> contain source code and run scripts for sample MPI applications, test programs and benchmarks.	<p>If you want to build them via the "Rebuild MPI Library and Tools" option in the <code>opafastfabric</code> menu, you must first perform the following steps to copy the files out of <code>/usr/lib</code> that may be a read-only directory:</p> <ol style="list-style-type: none"> 1) Create a mirror directory named <code>mpi_apps</code> in a writable area, such as under <code>\$HOME</code>. <pre>mkdir \$HOME/mpi_apps</pre> 2) Copy <code>/usr/lib/opa/src/mpi_apps/*</code> to the mirror directory. <pre>cp -r /usr/lib/opa/src/mpi_apps/* \$HOME/mpi_apps</pre> 3) Edit <code>/etc/sysconfig/opa/opafastfabric.conf</code> and change the setting of <code>FF_MPI_APPS_DIR</code> to the new mirror directory. <pre>export FF_MPI_APPS_DIR=\${FF_MPI_APPS_DI R:-\$HOME/mpi_apps}</pre> <p>You can now execute the "Rebuild MPI Library and Tools" option safely.</p>
135649	Software Installation/Packaging	The XPPSL kernel changes conflict with items in the SLES* 12 SP1 kernel RPM. This causes the recompile of the SLES* 12 SP1 <code>compat-rdma</code> package to have an error.	Intel does not recommend using SLES* 12 SP1 with the KNL-F for this release.
135729 135870	HFI Host Pre-boot Software	KNL-F/SKL-F ports are offline in pre-boot setting when connected with AOC.	If pre-boot connectivity is required, connect using QSFP cables or wait until OS load is complete to regain AOC connectivity.




Table 2-2. Open Issues (Sheet 5 of 6)

ID	Component	Description	Workaround
135812	Fabric Management Software	FM may crash and restart in the event of a failure during topology assignments. This may result in mismatched port physical states on a link. While unlikely, this event may occur when there are integrity issues on a link.	<p>1. Disable the PM completely. or</p> <p>2. Disable both sides of the offending link and restart FM. To get the link information, look for a log message in the following form:</p> <pre>fm0_sm[<PID>]: ERROR[<topology>]: SM: topology_assignments: Marking port [<PortNum>] of node [<#>] <NODE_DESC> guid <GUID> DOWN in the topology</pre> <p>This log message may occur about 10-12 Pm sweeps before the crash (usually 2 minutes).</p>
135929	HFI Host Pre-boot Software	Intel® Omni-Path Boot nodes occasionally dropped from fabric when switching master SM from one node to another.	Reboot PXE client node.
135958	HFI Host Fabric Software/ PSM2	Spurious segmentation faults with greater than 2MB PSM2 transfers on Intel® Xeon Phi™ platforms.	<p>Note: This workaround is needed only on Intel® Xeon Phi™ machines, because settings applied for Intel® Xeon® CPUs prevent this bug.</p> <p>The PSM2 environment variable PSM2_MQ_RNDV_HFI_WINDOW must be set to 2097152 or smaller (that is, 2MB maximum). In cases where the OS physical memory is fragmented beyond 512 physical regions, it may show a segfault attempting to transfer messages larger than 2MB. Limiting this value to 2MB or less prevents fragmentation getting beyond 512 physical regions and prevents the issue.</p> <p>Open MPI Example: (using -x)</p> <pre>#mpirun <job control> -x PSM2_MQ_RNDV_HFI_WINDOW=2097152 <other params></pre> <p>MVAPICH2 and Intel® MPI Example: (using -genv)</p> <pre>#mpirun <job control> -genv PSM2_MQ_RNDV_HFI_WINDOW=2097152 <other params></pre>
135975	Fabric Management Tools/FastFabric	After performing an OPA software configuration update, some unmanaged switches do not update the settings for LinkWidth and LinkWidthDnGrade enables.	A reboot is required for configuration changes made to an externally managed switch to become active.
136027	Fabric Management Tools/FastFabric	IFS hostverify.sh script does not provide reliable results for pstates_on and governor tests on RHEL* 7.3 and SLES* 12 SP2.	No workaround available.
136028	Fabric Management Tools/FastFabric	Two versions of the UEFI firmware are contained in the hfi-uefi RPM in the 10.3.0 IFS and BASIC packages. The files are functionally identical except the unsigned files (HfiPcieGen3Loader_<version number>.unsigned.rom and HfiPcieGen3_<version number>.unsigned.efi) are not signed for secure boot.	<p>Intel recommends using the signed file versions: HfiPcieGen3Loader_<version number>.rom and HfiPcieGen3_<version number>.efi</p> <p>The unsigned files may be used but will not support secure boot applications.</p>



Table 2-2. Open Issues (Sheet 6 of 6)

ID	Component	Description	Workaround
136152	Software Configuration Management	Server platforms using integrated HFI for OPA (commonly known as "-F") require BIOS that provides UEFI version 1.3.1.0.0 and a configuration data file for pre-boot support of Active Optical Cables (AOC). Some servers may not have these files available in BIOS and will therefore not support AOC in pre-boot.	Contact your server vendor for the appropriate BIOS update.
136160	HFI Host Hardware	On some Intel® Xeon Phi™ with integrated Intel® Omni-Path fabric platforms, the second integrated HFI is discovered first and is subsequently identified as the first HFI device. As a result, when issuing Intel® Omni-Path commands, the second HFI appears first in the results. In Linux* and various Intel® Omni-Path tools, the HFI reporting order may be the opposite of the order appearing on the Intel® Xeon Phi™ with integrated Intel® Omni-Path fabric cable/faceplate.	You can identify the second integrated HFI by inspecting the Node GUID or Port GUID/Port GUID reported by <code>opainfo</code> or other commands such as <code>hfi1_control -i</code> . Note that bit 39 of the PortGUID, the most significant bit, is set for the second HFI, and is clear for the first HFI. Keep in mind that when issuing various Intel® Omni-Path CLI commands targeted at a specific HFI using the <code>-h</code> option, <code>-h 1</code> correlates to the device that is listed as <code>hfi1_0</code> . As a result, the issued command affects the second HFI instance in cases where the second HFI port instance appears first.
136437	HFI Host Fabric Software	When using RHEL* 7.2, the default generic PXE boot image does not work due to missing driver and firmware files.	None.
136621	HFI Host Fabric Software	PCIe Fatal Errors during reboot cycles on server platforms using integrated HFI for OPA (commonly known as "-F").	If both WFR devices are enumerated and running at 8GT/s, you can ignore the error reporting and continue operations. Otherwise, reboot or power cycle the platform.
136628	Open Software	A bug in the Linux* kernel (CVE-2016-5195, also called Dirty COW) requires you to update the kernel for your operating system.	Contact Intel Customer Support to obtain a Technical Advisory containing detailed instructions for this process.
136723	Software Installation/Packaging	Upgrading your OPA installation from version 10.2 to 10.3 may not install the correct host driver.	Uninstall the 10.2 OPA software before installing 10.3 using the command: <code>./INSTALL -u</code> Install the 10.3 software using the command: <code>./INSTALL -a</code>

§ §